

(19)



Europäisches Patentamt
European Patent Office
Office européen des brevets

(11)



EP 1 229 430 A1

2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
07.08.2002 Bulletin 2002/32

(51) Int Cl.7: G06F 1/32, G06F 11/14

(21) Application number: 01410012.7

(22) Date of filing: 30.01.2001

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE TR

• Combe, Jean-Michel
38420 Domène (FR)

Designated Extension States:
AL LT LV MK RO SI

(74) Representative: Lloyd, Richard Graham
Intellectual Property Section,
Legal Department,
HEWLETT-PACKARD FRANCE,
Etablissement de Grenoble
38053 Grenoble Cedex 9 (FR)

(71) Applicant: Hewlett-Packard Company,
A Delaware Corporation
Palo Alto, CA 94304 (US)

(72) Inventors:

- Neuman, Paul
38640 Claix (FR)

(54) Power management system and method

(57) The present invention relates to a power management system and method for a computer system. The current ACPI specification defines, amongst other states, an S3 state and an S4 state. The S3 state has the advantage that a transition to a working state S0 is relatively quick, 5 seconds, but the disadvantage that the system context cannot be recovered after a power failure. The S4 state has the disadvantage that a trans-

sition to the S2 state takes a significant period of time, between 20-40 seconds, but the advantage that the system context can be recovered from a power failure. Suitably, the present invention defines a safe S3 state in which the transition to the working state takes a relatively short period of time but the same state, upon entry into that state, also saves system context data to a non-volatile storage device to allow recovery from a power failure.

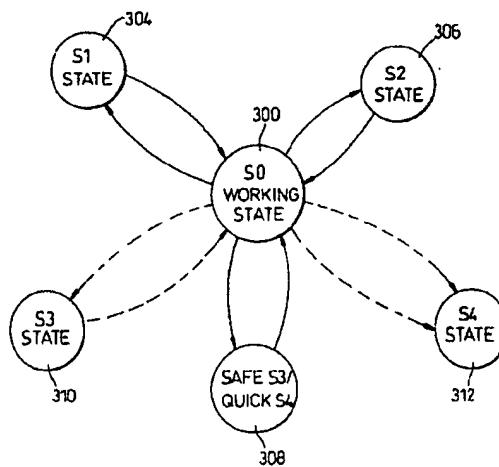


Fig. 3

EP 1 229 430 A1

Description

[0001] The present invention relates to a power management system and method, and, more particularly, to power management of computing systems.

[0002] The computing industry has developed a common interface for enabling robust operating system directed motherboard system configuration and power management (OSPM) of entire computer systems. The common interface definition and functionality manifests itself in the Advanced Configuration and Power Interface (ACPI) specification. The current version of the ACPI is version 2, having a release date of July 27, 2000 together with the ACPI Errata version 1.3, November 27, 2000, both of which are incorporated herein by reference for all purposes.

[0003] The ACPI specification defines a number of operating states for computer systems, such as, for example, desktop, mobile, workstation and server computers and laptop computers. Currently the ACPI specification defines five states, that is, states S0, S1, S2, S3 and S4. Each of the five states represents a different state or degree of power consumption of the associated computer system. State S0 represents the conventional operating state or system working state in which the computer system is fully functional and is not in a power saving mode. The remaining states, that is, states S1-S4, represent the system sleeping states in which the computer system has undertaken steps to at least reduce, and preferably minimise, power consumption.

[0004] Of particular interest are the S3 and S4 states. The ACPI specification defines the behaviour of the S3 state such that less power is consumed within the S3 state as compared to the S2 state; the processor does not execute instructions; the processor context is not maintained; the dynamic RAM context is maintained; power resources are held in a state that is compatible with the S3 state; devices associated with the computer system are operable such that they are compatible with the S3 state, that is, only devices that solely reference power resources are in the ON state (all other devices are in the off or D3 state); and devices that are enabled to wake the system, and that can do so from their current state, can initiate a hardware event that transitions the system to S0.

[0005] The ACPI specification defines the system behaviour in the S4 state as follows: the S4 state is logically lower than the S3 state and is arranged to consume less power than the S3 state; the processor(s) is(are) not executing instructions; the processor context is not maintained; RAM context is not maintained; power resources are in a state that is compatible with the S4 state, that is, all power resources that supply system level power in the S0, S1, S2 or S3 states are in the OFF state; and all devices are operable so as to be compatible with the current power resources states, that is, all devices are in the D3 state when the system is in the S4 state.

[0006] A system, upon entering the S3 state or in

preparation for entering the S3 state, saves, for example, the data necessary for resumption of the normal working state, S0, to RAM. Therefore, it can be appreciated that upon wake-up from the state S3, that the system context saved to RAM can be accessed and restored relative quickly. Therefore, the S3 sleeping state is known as a low wake-latency sleep state. However, the S3 state suffers from a major inadequacy in the event of a power failure that adversely affects the RAM such that the content of the RAM is lost. Such a power failure prevents a reliable transition to the working system state S0 and a re-boot of the computer system may be necessary. It can be appreciated that the data of applications and system context will be lost under such circumstances. Furthermore, a relatively long period of time will elapse during the re-boot before the computer system reaches the system working state S0.

[0007] In contrast, the OSPM of the system, before entering the S4 state, saves a significant amount of data to a non-volatile storage medium such as, for example, a HDD. Conventionally, data comprising the entire content of the RAM and the current system context are saved to a file, typically called "HIBERFIL.SYS", which is stored on the HDD.

[0008] Upon wake-up from the S4 state, the OSPM of the system is responsible for restoring the system context. Therefore, a system transition from the S4 state to the S0 state involves a significant amount of data recovery. The content of the file "HIBERFIL.SYS" is retrieved from the HDD and loaded into RAM and is used to restore the system context as it was at the time of entering the state S4. Due to the need to access a relatively slow storage device, access times of the storage device are relatively long. Therefore, the S4 state is considered to be the longest wake-latency sleeping state. Since a non-volatile storage medium is used to store the recovery data, the system state S4 will allow recovery from a power failure. However, the time taken to effect such a recovery is unacceptably long.

[0009] It is an object of the present invention to mitigate at least the problems of the prior art.

[0010] Accordingly, a first aspect of the present invention provides a method for power management of a system comprising a processor, a power supply and first storage medium and second non-volatile storage medium having first and second data access times respectively such that the first data access time is less than the second data access time; the system being operable in a plurality of states, each state having an associated level of system power consumption, the method comprising the steps of: saving data representing a current system context to allow a transition to a first state of the plurality of states from a second state of the plurality of states to both the first and second storage media; and placing the system in the second state; recovering the data from the first storage medium in response to detection of an event while the system is in the second state or, in the event that the power supply, while in the

second state, failed to supply sufficient power to the system to maintain the second state, recovering the data from the second storage medium; and restoring the system context using the recovered data.

[0011] Advantageously, embodiments of the present invention enable a power management system to be realised that provides, in the absence of a power failure, a relatively fast wake-up time from a sleep state and, following a power failure during such a sleep state, allows recovery of the system context to the sleep state.

[0012] A second aspect of the present invention provides a power management system comprising a processor, a power supply and first storage medium and second non-volatile storage medium having first and second data access times respectively such that the first data access time is less than the second data access time; the system being operable in a plurality of states, each state having an associated level of system power consumption, the system further comprising: means for saving data, representing a current system context to allow a transition to a first state of the plurality of states from a second state of the plurality of states, to both the first and second storage media; and means for placing the system in the second state; means for recovering the data from the first storage medium in response to detection of an event while the system is in the second state or, in the event that the power supply, while in the second state, failed to supply sufficient power to maintain the second state, recovering the data from the second storage medium; and means for restoring the system context using the recovered data.

[0013] Embodiments of the present invention will now be described by way of example only with reference to the accompanying drawings in which:

figure 1 shows schematically a computer system within which power management can be realised;

figure 2 illustrates schematically ACPI states and state transitions for a prior art power management system;

figure 3 depicts states and associated state transitions of a power management system according to a first embodiment;

figure 4 shows a flowchart of a power-off or sleep process according to an embodiment;

figure 5 depicts a flowchart of a wake-up or on process according to an embodiment; and

figure 6 illustrates a flowchart for system context recovery following a power failure during a sleep state.

[0014] Figure 1 illustrates schematically a computer system 100 within which ACPI specification power man-

agement can be realised. Figure 1 also illustrates other functional elements of the computer system 100. The computer system 100 comprises a processor 102, RAM 104, a non-volatile storage medium 106, such as, for example, a HDD. The motherboard (not shown) of the computer system comprises a LED 108 to provide a visual indication of the sleeping state of the computer system even when the system cover or housing has been removed. The motherboard also supports an ACPI BIOS 110 or at least a BIOS 110 which supports ACPI routines. As is conventional, the computer system has an operating system 112 which is arranged to implement operating system directed power management (OSPM) 114. The computer system will, of course, run various applications 116 and 118. Additional hardware and software functionality is provided in the form of power management event detection logic 120 which detects events in response to which the current state S0-S4 of the computer system should change to another state. For example, the user may depress an on button 124, in which case the computer system should effect a transition from a current sleeping state to a working state S0. Alternatively, the user may instigate a software shutdown of the computer system 100 in response to which the computer system should effect a transition from the current state S0 to a sleeping state. The events that the power management event detection logic 120 may detect also include other events, such as modem generated events signalling to the OSPM that data is being received and at least the modem and RAM should be suitably powered-up to allow reception of the data.

[0015] Referring to figure 2, there is shown a state diagram 200 for an OSPM system according to the prior art. The state diagram has five states, that is, states S0 202, S1 204, S2 206, S3 208 and S4 210. The five states are briefly described below.

[0016] State S0: While a system is in state S0, the system is said to be in a working state. The behaviour of that state is defined such that a processor 202, or, in a multi-processor system, the processors are, in one of a number of so-called processor states, C0 214, C1 216, C2 218, ..., CN 220, which each represent varying degrees of processor operation and associated power consumption; the dynamic RAM context is maintained by the processor; any devices 222 and 224 connected to, or forming part of, the system 100 are individually managed by the operating system software and can be in any one of four possible device states D0-D3, which, again, reflect varying degrees of power consumption; and any associated power resources are arranged to be in a state that is compatible with the device states.

[0017] State S1: The S1 state 204 is a low wake-up latency sleeping state. In this state, no system context is lost (CPU or chip set) and the system hardware maintains all system context.

[0018] State S2: The S2 state 206 is also considered to be a low wake-up latency sleeping state. The S2 state 206 is substantially similar to the S1 state 204 but for

the CPU and the system cache context being lost in the S2 state, since, typically, the operating system is responsible for maintaining cache and processor context. [0019] State S3: The S3 state 208 is a low wake-up latency sleeping state where all system context is lost except system memory. The CPU, cache and chip set context are lost in this state. However, the system hardware maintains memory context and restores some CPU and L2 configuration context. The S3 state 208 was described in detail above.

[0020] State S4: The S4 state 210 is the lowest power, longest wake-up latency sleeping state supported by ACPI. To reduce power consumption, preferably to a minimum, it is assumed that the hardware platform has powered-off all devices. Platform context is maintained. The S4 state 210 has been described in detail above.

[0021] Figure 3 shows a state transition diagram 300 for a power management system according to a first embodiment. Again, it can be seen that the state transition diagram 300 comprises at least the working system state S0 302. Preferably, the conventional states S0, S1 and S2 are also supported in the first embodiment. The states S0-S2 in the first embodiment are substantially identical in operation and realisation to the corresponding states described above in relation to figure 2 and in current ACPI specification.

[0022] Additionally, the state diagram 300 illustrates a new state, that is, a Safe S3/Quick S4 state 308 (SS3/QS4). The behaviour of the system in the SS3/QS4 can be characterised by the actions of saving substantially the same data to a non-volatile storage medium 106 as the conventional S4 state while concurrently maintaining the same data as the conventional S3 state in memory. Furthermore, in the SS3/QS4 state only the RAM remains in a powered state while all other aspects of the system adopt substantially the same powered state of the conventional S3 state but for the power management event detection logic 120 to allow a wake-up from that state. As discussed above, the data is saved in a file 126 that is called "HIBERFIL.SYS".

[0023] Therefore, if a power failure occurs while the system is in the SS3/QS4 state, the system context can be restored by loading the HIBERFIL.SYS file from, for example, the HDD 106 and restoring the system context. In contrast to the prior art power management state S3, if a power failure occurs, the system context at the time of power failure is recoverable.

[0024] In the absence of a power failure, the system context, when waking from the SS3/QS4 state, can be restored within a relatively short period of time, such as, for example, 5 seconds, that is, within a time scale that is comparable to the wake-up time for a conventional S3 state but with the additional security of also being recoverable from a power failure, unlike the conventional S3 state.

[0025] Preferably, once the context has been restored following a power failure, the system enters or resumes the SS3/QS4 state. However, it will be appreciated that

embodiments could be realised such that any one of the states could be entered upon recovery.

[0026] Referring to figure 4 there is shown schematically a flowchart 400 for an off process, that is, a process for notionally switching off the system that utilises the SS3/QS4 state. Upon detection of a power-off event by the power management event detection logic at step 402, the OSPM saves the entire system context to the HDD 106 in the above mentioned file at step 404. In step 406, the power to the internal motherboard LED 108, which provides a visual indication of the sleep state of the system in the event that the system housing has been opened, is removed and the system hardware is arranged to adopt the same power saving configuration as the conventional S3 state in step 408. Once step 404 has been performed, the system has power for the RAM and the recovery data, or S4 data 126, has been stored on the HDD 106.

[0027] Referring to figure 5, there is shown a flowchart 500 for an "on" process which uses the SS3/QS4 state according to an embodiment. At step 502, the power management event detection logic 120 detects or receives a wake-up event and informs the wake-up and sleep logic 122 of the event. The wake-up event causes the wake-up and sleep logic 122 to instruct the OSPM to restore the system context from RAM at step 504. It can be appreciated that the wake-up process is substantially the same as the conventional S3 wake-up process.

[0028] Figure 6 shows a flowchart 600 of a restore process to recover from a power failure while the system was in the SS3/QS4 state. A power return event is detected by the power management event detection logic 120 at step 602 which causes the system to recover the previously stored system context data 126 from, for example, the HDD 106 and to load that data into RAM 104 at step 604. The power configuration of the system is arranged by the OSPM 114 to adopt substantially the same power configuration as in the conventional S3 state in step 606.

[0029] A transition from the conventional S3 state to the working state, that is, state S0, takes approximately 5 seconds as does the transition to the S0 state from the SS3/QS4 state, which are both significantly quicker than the current 20-40 second wake-up time for an S4 to S0 transition. However, the SS3/QS4 state has the additional advantage of allowing a consistent or safe recovery from a power failure while the system was in the power saving state SS3/QS4.

[0030] Although the above embodiments show a lack of support for S3 and S4 states, it will be appreciated that embodiments can be realised in which the S3 and S4 states are supported in addition to the states of the above embodiments. The states S3 310 and S4 312 are shown in figure 3 as being optionally supported

[0031] Furthermore, even though the above embodiments have been described in terms of having a number of system states, the present invention is not limited to

such system states. An embodiments can be realised in which other states such as, for example, Legacy states, mechanical-off states G3 and soft-off S5 states are also supported.

[0032] Although the above embodiments use a HDD as the non-volatile storage means, it will be appreciated that other forms of non-volatile storage may be used. For example, a flash-memory may be used to store the data to allow recovery from a power failure.

Claims

1. A method for power management of system comprising a processor, a power supply and a first storage medium and a second non-volatile storage medium having first and second data access times respectively such that the first data access time is less than the second data access time; the system being operable in a plurality of states, each state having an associated level of system power consumption, the method comprising the steps of: saving data representing a current system context to allow a transition to a first state of the plurality of states from a second state of the plurality of states to both the first and second storage media; and placing the system in the second state; recovering the data from the first storage medium in response to detection of an event while the system is in the second state or, in the event that the power supply, while in the second state, failed to supply sufficient power to the system to maintain the second state, recovering the data from the second storage medium; and restoring the system context using the recovered data.
2. A method as claimed in claim 1 further comprising the step of entering, after restoring the system context, a selectable one of the plurality of states.
3. A method as claimed in either of claims 1 and 2, further comprising the step of entering, after restoring the system context, the second state of the plurality of states.
4. A method as claimed in either of claims 1 and 2, further comprising the step of entering, after restoring the system context, the first state of the plurality of states.
5. A method as claimed in any preceding claim in which the step of saving data representing the current system context comprises saving at least one of data relating to a current operational state of the processor and data relating to the content of the first storage medium.
6. A method as claimed in any preceding claim in which the current operational state of the processor

is one of a plurality of operational states of the processor.

7. A method as claimed in any preceding claim in which the system further comprises at least one device, and in which the step of saving data relating to the current system context comprises saving data relating to a current operational state of the at least one device.
8. A method as claimed in claim 7 in which the current operational state of the device is one of a plurality of operational states of the device.
9. A power management system comprising a processor, a power supply and a first storage medium and a second, non-volatile, storage medium having first and second data access times respectively such that the first data access time is less than the second data access time; the system being operable in a plurality of states, each state having an associated level of power consumption, the system further comprising: means for saving data representing a current system context to allow a transition to a first state of the plurality of states from a second state of the plurality of states to both the first and second storage media; and means for placing the system in the second state; means for recovering the data from the first storage medium in response to detection of an event while the system is in the second state or, in the event that the power supply, while in the second state, failed to supply sufficient power to maintain the second state, recovering the data from the second storage medium; and means for restoring the system context using the recovered data.
10. A system as claimed in claim 9 further comprising means for entering, after restoring the system context, a selectable one of the plurality of states.
11. A system as claimed in either of claims 9 and 10, further comprising means for entering, after restoring the system context, the second state of the plurality of states.
12. A system as claimed in either of claims 9 and 10, further comprising means for entering, after restoring the system context, the first state of the plurality of states.
13. A system as claimed in any of claims 9 to 12 in which the means for saving data representing the current system context comprises means to save at least one of data relating to a current operational state of the processor and data relating to the content of the first storage medium.

14. A system as claimed in any of claims 9 to 13 in which the current operational state of the processor is one of a plurality of operational states of the processor.
15. A system as claimed in any of claims 9 to 14 in which the system further comprises at least one device and means for saving data relating to a current operational state of the at least one device. 5
16. A system as claimed in claim 15 in which the current operational state of the device is one of a plurality of operational states of the device. 10
17. A computer program element for implementing a method as claimed in any of claims 1 to 8 or a system as claimed in any of claims 9 to 16. 15
18. A computer program product comprising a computer readable storage medium having stored thereon a computer program element as claimed in claim 17. 20

25

30

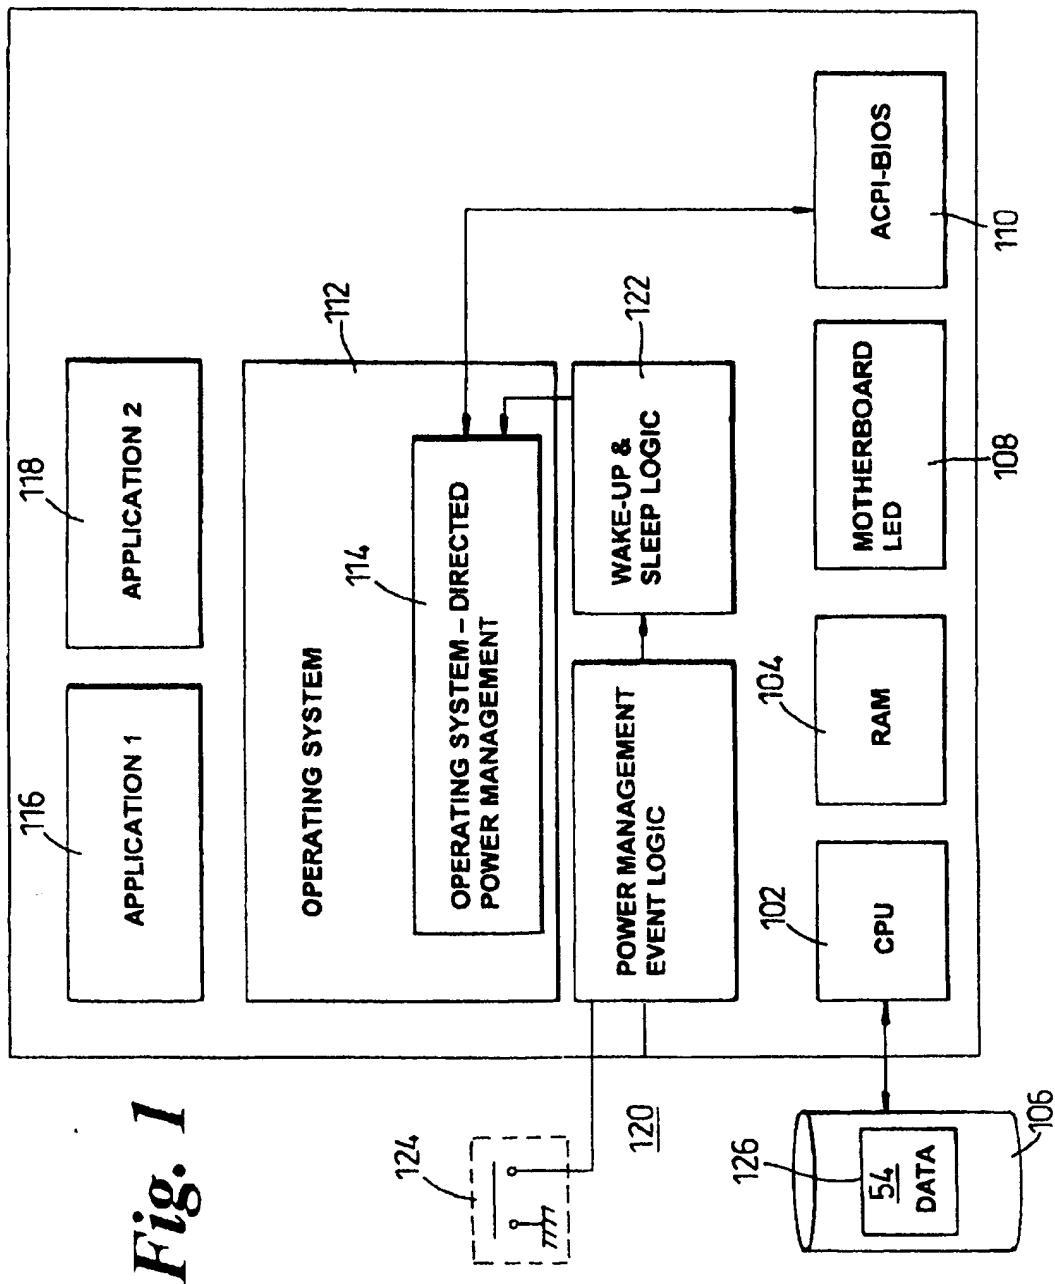
35

40

45

50

55

100*Fig. 1*

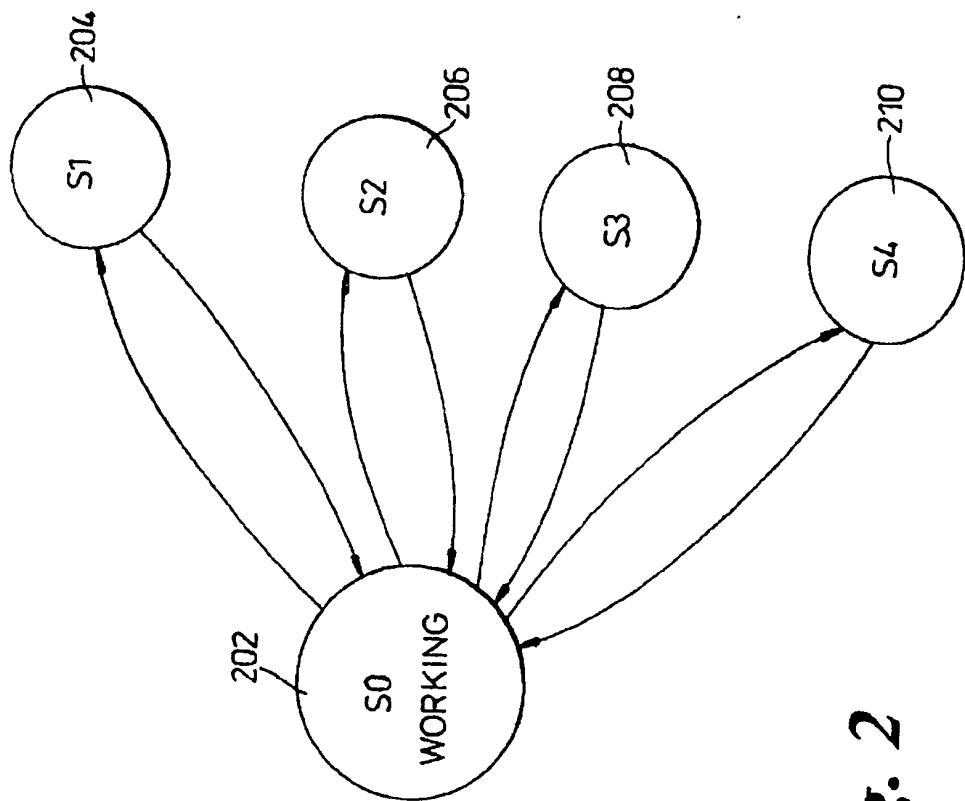
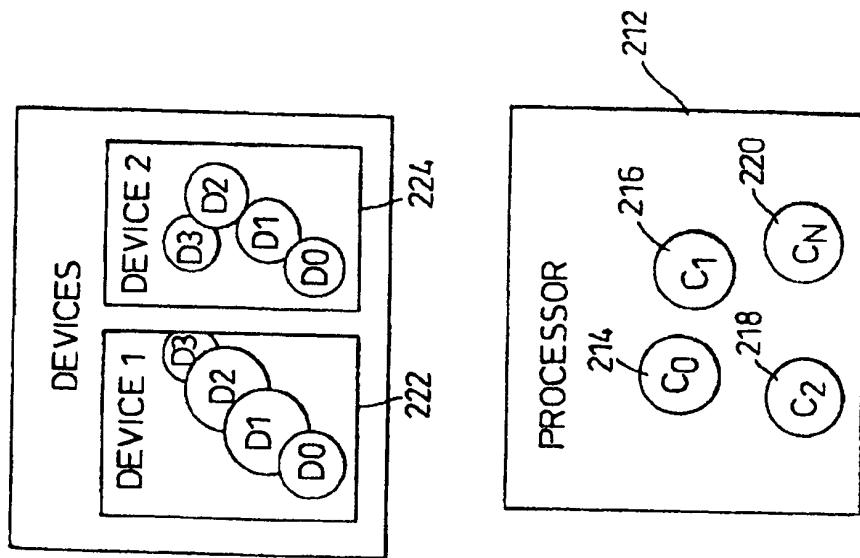


Fig. 2



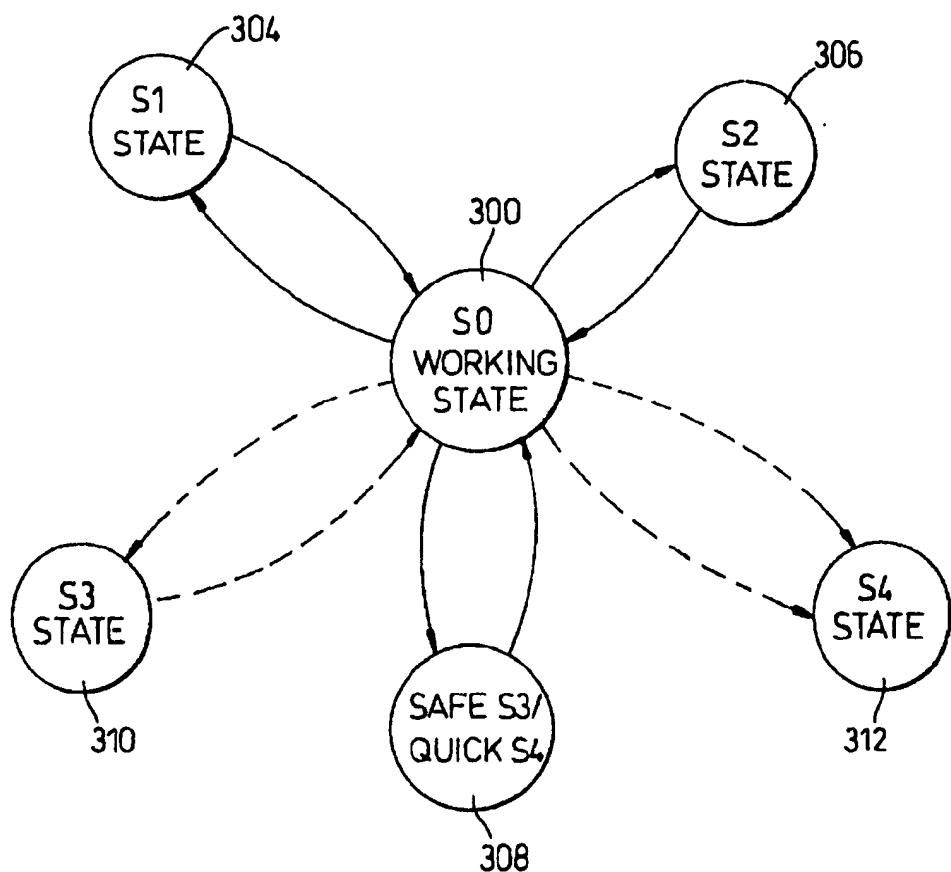


Fig. 3

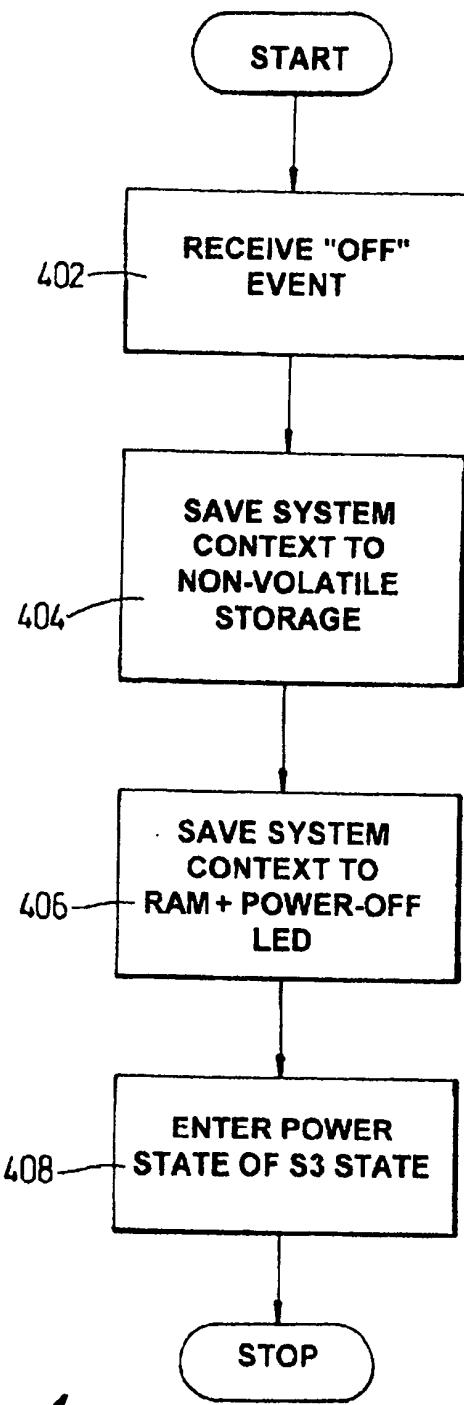


Fig. 4

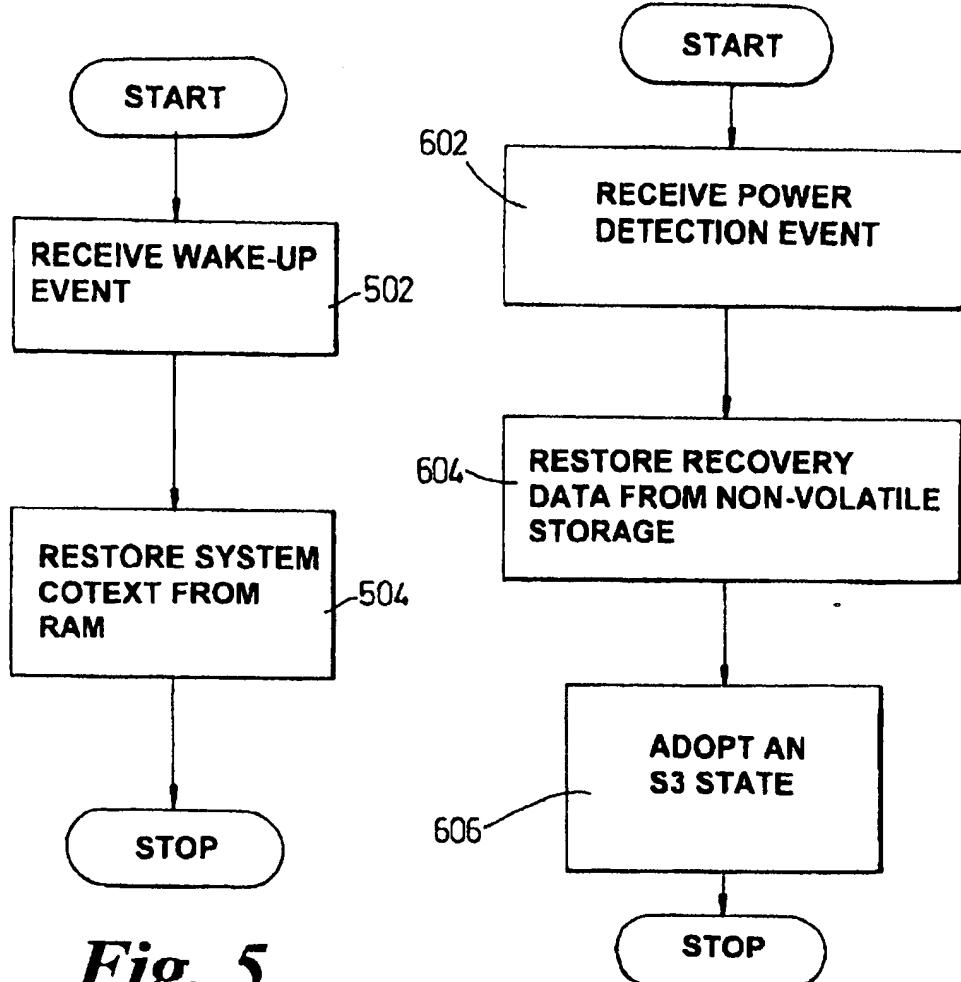


Fig. 5

Fig. 6



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 01 41 0012

DOCUMENTS CONSIDERED TO BE RELEVANT		
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim
X	US 5 978 922 A (ARAI ET AL.) 2 November 1999 (1999-11-02) * column 1, line 11 - line 57 * * column 2, line 14 - line 27 * * column 3, line 26 - line 38 * * column 5, line 7 - line 56 * * column 6, line 65 - column 8, line 47; figures 1,4,5 *	1,4-9, 12-18
A	US 6 131 166 A (WONG-INSLEY) 10 October 2000 (2000-10-10) * column 9, line 20 - column 14, line 37; figures 3-6 *	1,5,6,9, 13,14,17
The present search report has been drawn up for all claims		
Place of search	Date of completion of the search	Examiner
MUNICH	26 July 2001	Sepple, M
CATEGORY OF CITED DOCUMENTS		
X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document		
T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document		

ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.

EP 01 41 0012

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report. The members are as contained in the European Patent Office EDP file on The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

26-07-2001

Patent document cited in search report		Publication date		Patent family member(s)		Publication date
US 5978922	A	02-11-1999		JP 2988866 B JP 9237128 A EP 0793175 A		13-12-1999 09-09-1997 03-09-1997
US 6131166	A	10-10-2000		NONE		

EPO FORM P0459

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82

THIS PAGE BLANK (USPTO)

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- BLACK BORDERS**
- IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- FADED TEXT OR DRAWING**
- BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- SKEWED/SLANTED IMAGES**
- COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- GRAY SCALE DOCUMENTS**
- LINES OR MARKS ON ORIGINAL DOCUMENT**
- REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- OTHER:** _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.

THIS PAGE BLANK (USPTO)